arXiv:1603.04628v1 [physics.flu-dyn] 15 Mar 2016

# Accelerating a hybrid continuum-atomistic fluidic model with on-the-fly machine learning

David Stephenson[a,*], James R. Kermode[b], Duncan A. Lockerby[a]

[a] *School of Engineering, University of Warwick, Coventry CV4 7AL, UK*
[b] *Warwick Centre for Predictive Modelling, School of Engineering, University of Warwick, Coventry CV4 7AL, UK*

**Abstract**

We present a hybrid continuum-atomistic scheme which combines molecular dynamics (MD) simulations with on-the-fly machine learning techniques for the accurate and efficient prediction of multiscale fluidic systems. By using a Gaussian process as a surrogate model for the computationally expensive MD simulations, we use Bayesian inference to predict the system behaviour at the atomistic scale, purely by consideration of the macroscopic inputs and outputs. Whenever the uncertainty of this prediction is greater than a predetermined acceptable threshold, a new MD simulation is performed to continually augment the database, which is never required to be complete. This provides a substantial enhancement to the current generation of hybrid methods, which often require many similar atomistic simulations to be performed, discarding information after it is used once.

We apply our hybrid scheme to nano-confined unsteady flow through a high-aspect-ratio converging-diverging channel, and make comparisons between the new scheme and full MD simulations for a range of uncertainty thresholds and initial databases. For low thresholds, our hybrid solution is highly accurate — within the thermal noise of a full MD simulation. As the uncertainty threshold is raised, the accuracy of our scheme decreases and the computational speed-up increases (relative to a full MD simulation), enabling the compromise between precision and efficiency to be tuned. The speed-up of our hybrid solution ranges from an order of magnitude, with no initial database, to cases where an extensive initial database ensures no new MD

*Corresponding author
*Email address:* david.stephenson@warwick.ac.uk (David Stephenson )

simulations are required.

*Keywords:* Multiscale modelling, Machine Learning, Hybrid methods, Micro/nanofluidics, Molecular Dynamics

## 1. Introduction

Almost all fluid engineering systems are multiscale in their nature. At the smallest scale, the fluid and surrounding environment are comprised of atoms, with interactions occurring across nanometers ($10^{-9}$ m) and over femtoseconds ($10^{-15}$ s), while the fluid flow is characterised by the scale of the system geometry, which is often many orders of magnitude larger. In most instances, the separation of scales is so large that the atomistic behaviour can be accurately incorporated into a continuum fluid description through empirical boundary conditions (e.g. no-slip at the walls) and constitutive relations (e.g. viscosity in the shear stress – strain rate relation). However, as some characteristic dimension of the system approaches the nanoscale (and microscale in gases), these approximations break down, and the fluid flow becomes highly dependent on atomistic phenomena [1–4].

There are numerous applications in fluid dynamics where atomistic information is required to capture non-continuum/non-equilibrium phenomena, but the macroscopic flow develops over much larger length and time scales: for example, pumping technology that exploits thermal creep in a rarefied gas [5]; and low-power, high-throughput nanotube membranes for salt water desalination [6]. The multiscale nature of these systems leads to a dual requirement for capturing the local atomistic-scale interactions and macroscale fluid response. The complexity of the flow necessitates modelling with atomistic resolution, but the state-of-the-art techniques (molecular dynamics (MD) for dense fluid flows [7], and the direct simulation Monte-Carlo method (DSMC) for rarefied gas flows [8]) are extremely computationally expensive. This limits their application to small system sizes, typically $\mathcal{O}(100\ \text{nm}^3)$, and short simulation times, typically $\mathcal{O}(100\ \text{ns})$, rendering many important engineering problems intractable, and limiting possibilities for comparison with experiments.

Hybrid methods provide a promising framework for simulating such systems by combining continuum (macro) and atomistic (micro) solvers and exploiting scale-separation where it exists to obtain a high accuracy, yet efficient, solution. There has already been extensive research into hybrid

approaches in fluid dynamics (see, for example, recent reviews [9–12]), with most methodological variations concerning either a) the coupling strategy between the macro and micro models, or b) the boundary conditions that are imposed on each model. Coupling is usually performed via the exchange of state variable or flux information, to ensure consistency between the macro and micro models at locations where an atomic resolution is required. The information passed from one model is then applied as a boundary condition to locally constrain the other model.

Broadly speaking, the types of problems that hybrid methods are designed to model can be divided into three categories according to common features: Type A, Type B, or Type C. In Type A problems, a continuum model provides a reasonably accurate solution for much of the flow field, and an atomistic solver is only required for regions that contain non-continuum/non-equilibrium effects such as an isolated defect, a singularity, or regions near nanoscale surfaces. This is often referred to as domain decomposition because the macro and micro domains are coupled across spatial interfaces. Examples of Type A problems include the moving contact line problem [2] and resolving singularities that occur in the corners of lid-driven cavity flows [13]. In Type B problems, a continuum model spans the entire system domain, and an atomistic resolution is required everywhere to continually refine the continuum solution by providing local constitutive/boundary information. This is an embedded coupling approach, with the macro and micro domains occupying the same physical space; an example of a Type B problem is capturing the non-Newtonian flow of a complex nanofluid [14]. In Type C problems, high-aspect-ratio flows result in non-continuum/non-equilibrium effects persisting over the entire (or majority of the) cross section of the fluid stream, while hydrodynamic properties vary slowly in the streamwise direction. The former must be captured by an atomistic solver, while the latter can be described by a continuum. This is also an embedded coupling approach, with examples including the flow through nanotube desalination membranes [6] and air-layer lubrication of liquid journal bearings [15]. However, despite much progress, a common flaw of hybrid methods is that, often, similar and repetitious atomistic simulations are performed — i.e. information is used once, then wastefully discarded, despite configuration conditions remaining similar.

In recent years, machine learning (ML) techniques have been used in an increasingly wide range of disciplines [16]. Prominent examples based on Bayesian inference include the calibration of computer models [17] or the

closely related problem of developing surrogate models that can be computed much more cheaply [18]. These approaches can help to quantify uncertainty; for example, the uncertainty induced in density functional theory (DFT) calculations from the choice of exchange-correlation functional [19], or the uncertainty induced in nanoscale flow simulations from the choice of interatomic potentials [20]. Hand-calibrated force fields have been successfully replaced by models built automatically from quantum mechanical (QM) reference data, using neural networks [21] or Gaussian process regression [22]. This has enabled accurate models to be produced for materials that were conventionally difficult to describe precisely with interatomic potentials [23]. However, all such approaches are limited by the training data used to fit the ML model. In particular, they are not transferable to situations not envisaged at the time of construction.

More recently, a new class of methods for ML-based atomistic simulation have been proposed, which addresses this limitation by avoiding ever attempting to build a full description of the reference database (i.e. potential energy surface). Instead, they accelerate MD simulations by computing new reference data 'on the fly' where and when it is needed [24, 25]. The reference data is added to a growing database that is used to infer the information needed to propagate the dynamics, greatly reducing the computational cost of MD simulations while retaining the link to QM-accurate methods [26]. Similar ideas have been applied to develop neural network models for hybrid continuum-atomistic flow simulations [27]. However, optimising the architecture of parametric neural networks, while avoiding over- or under-fitting, is a challenging task. Moreover, in [27], the criterion for computing new reference data is somewhat ad hoc, based on a fixed difference between parameters in input space. Here, we improve upon this idea by adopting a non-parametric approach to regression — using a Gaussian process to model the atomistic database in a hybrid continuum-atomistic scheme. The non-parametric model helps avoid over- or under-fitting to the data, and employing Bayesian inference for output predictions means that the criterion for obtaining new reference data can be based on a target threshold variance in output space, leading to a new robust approach to hybrid modelling in fluid dynamics, with close-to-optimal information efficiency.

## 2. Methodology

The specifics of our new scheme will depend on the exact hybrid method chosen, but the basic idea remains the same. Each time the continuum solver seeks to obtain atomistic information, rather than immediately performing a new simulation, we first predict the simulation output based on the similarity between the current system configuration and all configurations in a pre-existing database. A new atomistic simulation is only performed when the uncertainty of this prediction is above a predetermined threshold, i.e. the system configuration is not well represented in the database; otherwise the prediction is deemed sufficient. The data from new simulations is used to augment the existing database, whose completeness is never required. As this database grows, progressively fewer atomistic simulations are required, because 'new' system configurations are more rarely encountered.

The system we will use as a benchmark for our new scheme is dense fluid flow through a converging-diverging channel, with the flow driven by a time-variant periodic external force $F_{\text{ext}}(t)$. The geometry of the system is presented in Fig. 1a. We choose this as our benchmark system for two main reasons: 1) it is multiscale both spatially and temporally, demonstrating the potential for high computational savings; and 2) the results for the full atomistic simulation, recently published by Borg et. al. [28], provide a useful basis for comparison. The hybrid method we use to model this system was developed by Borg et. al. [28], and described in detail therein. However, for completeness, we provide a short summary below.

### *2.1. Hybrid method*

The benchmark system is a Type C problem, so non-continuum effects persist over the entire cross section, and spatial scale separation can only be exploited in the streamwise direction. As such, atomistic subdomain simulations must cover the entire channel height, and are placed at regular intervals in the streamwise ($s$-) direction according to the equation:

$$s_i = \frac{(i-1)}{N} L, \tag{1}$$

where $i$ is the subdomain index, $N$ is the number of subdomains, and $L$ is the channel length. Note, the channel uses periodic boundary conditions (PBCs) in the ($s$-) direction, so the first subdomain is simultaneously located at the inlet and the outlet. The number of subdomains is set large enough to resolve

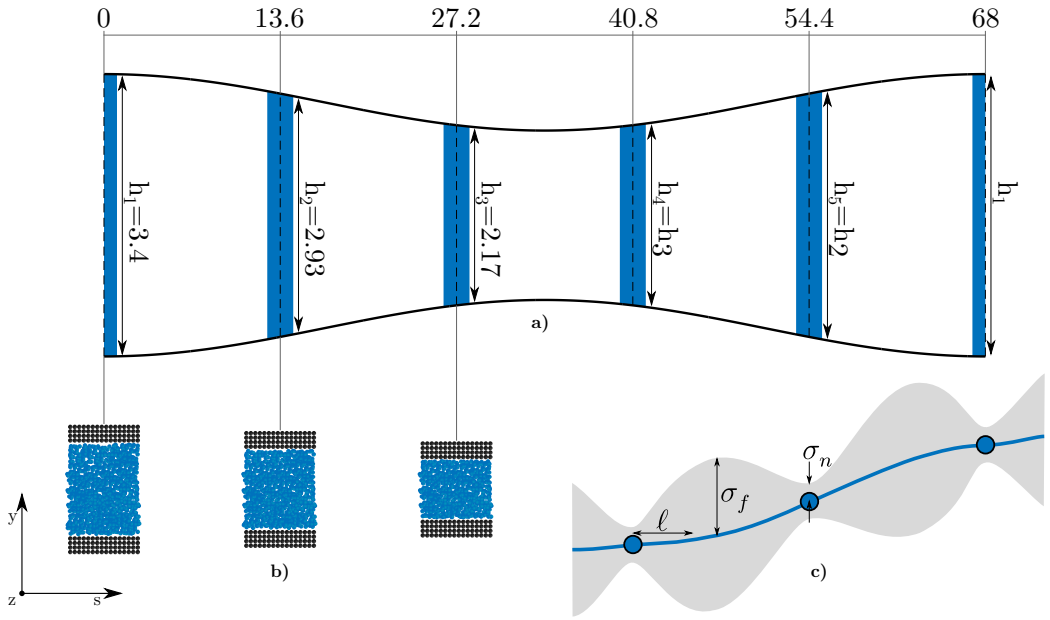


Figure 1: Schematics of a) the multiscaled converging-diverging nanochannel, b) the micro subdomain decomposition, and c) the representation of the Gaussian-process predictive model, where the blue line is the mean and the grey envelope represents one standard deviation. All dimensions are in nm.

the streamwise geometrical variation; here we, like Borg et. al., use $N = 5$. The channel height $h(s)$ is a function of the streamwise position $s$:

$$h = a\left[\cos\left(\frac{2\pi s}{L}\right) - 1\right] + h_1, \tag{2}$$

where $a = 0.68$ nm is a constant and $h_1 = 3.4$ nm is the channel height at the inlet/outlet (see Fig. 1). Temporal scale separation exists because the characteristic time for the evolution of the macro model (e.g. the period of the external force) is much larger than the characteristic time for the development of the micro model (e.g. the start-up time from rest); therefore, each micro subdomain simulation is considered to be in a quasi-steady state.

The macro model consists of the unsteady one-dimensional equations for mass and momentum conservation. We use MD for the micro model, with atoms interacting through pairwise potentials and moving according to Newton's laws of motion (see Appendix A for details). Coupling is performed by ensuring that the mass and momentum in each MD subdomain are consistent with the conservation laws of the macro model. For mass, we obtain

$$\frac{\partial \rho}{\partial t} + \left(\frac{1}{A}\right)\frac{\partial q}{\partial s} = 0, \tag{3}$$

where $\rho(s,t)$ is the density, $A$ is the cross-sectional area, and $q(s,t)$ is the mass flow rate which are time-averaged measurements in each MD subdomain simulation. For simplicity, the MD subdomains are simulated using PBCs in the $s$-direction (see Fig. 1b), which cannot support a pressure gradient. Therefore, for momentum flux to be hydrodynamically equivalent to that in the macro model, the total (external) force $F(s,t)$ applied to each atom is

$$F = F_{\text{ext}} - \left(\frac{m}{\rho}\right)\frac{\mathrm{d}p}{\mathrm{d}s}, \tag{4}$$

where $m$ is the mass of a single atom.

*2.2. Machine learning*

For the machine learning model, we use Gaussian process (GP) regression. Here we provide only a brief overview of the approach, see Rasmussen and Williams [18] for further details. A non-parametric approach to regression is to model the unknown underlying function $f(\mathbf{x})$ as random, with a GP distribution:

$$f(\mathbf{x}) \sim \mathcal{GP}\big(\mu(\mathbf{x}), K(\mathbf{x}, \mathbf{x}')\big), \tag{5}$$

where $\mathbf{x}$ is a point in input space, $\mu(\mathbf{x})$ is the mean function and $K(\mathbf{x}, \mathbf{x}')$ is the covariance function for the process of interest, which models the covariances between function outputs $y$ and $y'$ at $\mathbf{x}$ and $\mathbf{x}'$, respectively. For our system, the input vector $\mathbf{x} = \{h, \rho, F\}$ is composed of three macroscopic input variables: channel height $h$, density $\rho$, and force $F$. Here, we will use two separate GPs to predict the output for the mass flow rate $y = q$ and pressure[1] $y = p$ (note, for pressure, the input vector only depends on channel height and density).

A GP is a generalisation of a Gaussian probability distribution to an infinite number of dimensions. For each point in input space $\mathbf{x}$, we model the output $y = f(\mathbf{x})$ as a random variable with a univariate Gaussian distribution; similarly, we can model the output across the infinite range of input space using a GP. Starting from some prior belief in our function $P(f)$ (defined by our choice of covariance function, discussed below), we perform $M$ training MD simulations with inputs

$$X = \begin{bmatrix} h_1 & h_2 & \dots & h_M \\ \rho_1 & \rho_2 & \dots & \rho_M \\ F_1 & F_2 & \dots & F_M \end{bmatrix}, \tag{6}$$

and observe outputs $\langle \mathbf{y} \rangle$ (specifically, mass flow rate $\langle \mathbf{q} \rangle$ and pressure $\langle \mathbf{p} \rangle$). Using Bayesian inference, this leads to a posterior belief in the function $P(f|\langle \mathbf{y} \rangle, X)$ given by

$$P(f|\langle \mathbf{y} \rangle, X) = \frac{P(\langle \mathbf{y} \rangle | X, f) P(f)}{P(\langle \mathbf{y} \rangle | X)}, \tag{7}$$

where $P(\langle \mathbf{y} \rangle | X, f)$ is the likelihood of observing the outputs, given the prior belief, and $P(\langle \mathbf{y} \rangle | X)$ is the marginal likelihood. The angle brackets represent time-averaged measurements from MD simulations. The mean of the posterior $\mathbb{E}[f|\langle \mathbf{y} \rangle, X]$ will closely resemble the measured outputs $\langle \mathbf{y} \rangle$ near the input data points $X$, and the uncertainty (variance) will be low; however, the uncertainty away from the observed data remains large (see Fig. 1c). The posterior predictions thus become more accurate as the database grows and covers more of input space. As each data point 'speaks', the GP can be considered to have a finite, but unbounded, number of parameters, which grow with the database.

[1] averaged across the channel cross section.

To model the covariance between outputs, we use the squared exponential kernel

$$K(\mathbf{x}_i, \mathbf{x}_j) = \sigma_f^2 \exp\left(-\frac{d_{ij}^2}{2\ell^2}\right), \tag{8}$$

where $\sigma_f$ and $\ell$ are hyperparameters — with $\sigma_f^2$ representing the variance of the unknown function and $\ell^2$ representing the length scale[2] — and $d_{ij}^2$ is the normalised Euclidean distance between the points $\mathbf{x}_i = (h_i, \rho_i, F_i)$ and $\mathbf{x}_j = (h_j, \rho_j, F_j)$ in input space:

$$d_{ij}^2 = \left(\frac{h_i - h_j}{\overline{\Delta h}}\right)^2 + \left(\frac{\rho_i - \rho_j}{\overline{\Delta \rho}}\right)^2 + \left(\frac{F_i - F_j}{\overline{\Delta F}}\right)^2. \tag{9}$$

The denominator of each term in equation (9) is the mean separation between inputs in the database. This normalisation ensures that each input variable has approximately equal weight, and is recalculated when the covariance matrix is updated as the database grows. For the pressure kernel, only the first two terms in equation (9) are included.

When making predictions for the outputs, we include the inherent noise in the measured outputs from MD simulations by adding Gaussian noise with variance $\sigma_n^2$ to the covariance matrix

$$C(X, X) = K(X, X) + \sigma_n^2 I, \tag{10}$$

where $I$ is the identity matrix and $\sigma_n^2$ can be interpreted as an additional hyperparameter (see §2.3).

To make predictions for a set of test configurations with inputs $X_*$, it can be shown that Eq. 7 leads to a posterior distribution for the outputs $\mathbf{y}_*$ given by

$$\mathbf{y}_* | X, \mathbf{y}, X_* \sim \mathcal{N}\left(\bar{\mathbf{y}}_*, \mathrm{cov}(\mathbf{y}_*)\right), \tag{11}$$

where the mean is

$$\bar{\mathbf{y}}_* = \mu(X_*) + K(X_*, X) C(X, X)^{-1} \left(\langle \mathbf{y} \rangle - \mu(X)\right), \tag{12}$$

and the covariance matrix

$$\mathrm{cov}(\mathbf{y}_*) = K(X_*, X_*) - K(X_*, X) C(X, X)^{-1} K(X, X_*). \tag{13}$$

[2]Note, this is a length scale in input space of distances $d_{ij}$, and should not be confused with any physical length scale used in the multiscale modelling.

The variances $\boldsymbol{\sigma}^2_{y_*}$ of the new predictions are the diagonal of the posterior covariance matrix $\mathrm{cov}(\mathbf{y}_*)$. In our scheme, whenever the predicted standard deviation $\sigma_{y_*}$ exceeds a pre-determined uncertainty threshold $\sigma_t$, the prediction is deemed insufficiently accurate and a new MD simulation is automatically performed and added to the teaching database.

### *2.3. Hyperparameters*

Our non-parametric model nevertheless includes three hyperparameters $\boldsymbol{\theta} = (\sigma_f^2, \ell^2, \sigma_n^2)$, where $\sigma_f^2$ is the signal variance, i.e. the expected variance of the output function away from any known data points; $\sigma_n^2$ is the noise variance, i.e. the variance at known data points (see Fig. 1); and $\ell^2$ is the length scale, i.e. a measure of the distance between input points beyond which the outputs become uncorrelated; see Fig. 1c for a visual representation. The normalisation used here, in equation (9), makes the choice of hyperparameters highly intuitive; it is likely that $\sigma_f^2$ and $\ell^2$ will be approximately unity, and $\sigma_n^2$ will be some fraction of $\sigma_f^2$. We verified this intuition by computing the marginal likelihood of observing the measured outputs $\langle \mathbf{y} \rangle$ given the inputs $X$ and hyperparameters $\boldsymbol{\theta}$, given by

$$\log P(\langle \mathbf{y} \rangle | X, \boldsymbol{\theta}) = -\frac{1}{2} \langle \mathbf{y}^\top \rangle C(X,X)^{-1} \langle \mathbf{y} \rangle - \frac{1}{2} \log |C(X,X)| - \frac{n}{2} \log 2\pi \quad (14)$$

where it should be noted $C(X,X)$ depends on $\boldsymbol{\theta}$. We wish to maximise $\log P(\langle \mathbf{y} \rangle | X, \boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$, i.e. find the hyperparameters which maximise the likelihood of the outputs we observed. We calculated the log marginal likelihood over a small sample of training data points, which would later be used for initial databases (see Case D5 from Table 1 in §3.2). We did not optimise for $\sigma_n^2$ as this could be estimated directly from the data, by calculating the standard deviation of the mean output for each MD simulation (using the raw time-instantaneous data) and averaging over the number of MD simulations. We found $\sigma_{nq} = 0.05$ ng/s for mass flow rate and $\sigma_{np} = 0.003$ MPa for pressure[3]. Solving 14 numerically to maximise the log marginal likelihood for the other two hyperparameters, we obtained $\sigma_f = 1$ (ng/s and MPa for mass flow rate and pressure, respectively) and $\ell = 1$ for both mass flow rate and pressure, confirming our intuition.

[3]Note, we present the hyperparameters in terms of standard deviation, rather than variance, for the sake of clarity, as they share units with the outputs.

*2.4. Implementation*

The procedure for our machine-learning-accelerated hybrid method is described in Algorithm 1. After the initial databases for pressure and mass flow rate are loaded, and the number of subdomains $N$ is chosen, the location $\mathbf{s}_{1...N}$ and height $\mathbf{h}_{1...N}$ of the subdomains are calculated using equations (1) and (2), respectively. The subdomains are initialised with densities $\boldsymbol{\rho}^1_{1...N} = \{1331, 1320, 1278, 1273, 1312\}$ kg/m$^3$, equivalent to the initial densities in the full MD simulation.

The nature of the hybrid method means the pressure needs to be evaluated before the mass flow rate, as the mass flow rate is dependent on the force, which is in turn dependent on the pressure gradient. This is problematic because, when we encounter a 'new' configuration (i.e. one for which we cannot accurately predict the output), we do not want to perform two separate MD simulations for pressure and mass flow rate. We overcome this problem by ensuring the pressure prediction is always accurate, i.e. the pressure database is well represented across all of input space. This is more simple than it first sounds, because pressure is only dependent on two inputs: channel height and density. The channel heights at each subdomain are known and unchanging throughout a single case, so to robustly map density we only need to perform one pre-simulation per subdomain. In each pre-simulation, we start from a high density of 1800 kg/m$^3$ and measure the pressure as atoms are progressively removed, down to a density of 800 kg/m$^3$. Our GP regression for pressure is therefore similar to the empirical equation of state of Borg et. al. [28]. However, using a GP, we avoid having to guess the functional form of this equation of state. These pre-simulations are sufficiently cheap that they do not affect our overall computational efficiency calculations.

As the macro model evolves in time, a configuration $\mathbf{x}_* = \left\{h^i_j, \rho^i_j, F^i_j\right\}$ is tested to provide micro information, where $j = 1, N$ is the spatial index (representing subdomain location) and $i = 1, t_{\text{end}}/\Delta t$ is the temporal index; $t_{\text{end}}$ and $\Delta t$ are the end time and time step, respectively, for the macro model. If ${\sigma_q}^i_j > \sigma_t$, then a new MD simulation is performed with input $\mathbf{x}_*$. In our MD subdomain simulations, output measurements are averaged over 40000 time-steps (see Appendix A, MD simulation details) after the initial start-up phase. To minimise the start-up phase, new MD simulations are initiated with the final atomistic positions and velocities of a previously simulated configuration from our database; the input for this prior configuration $\mathbf{x}_{\text{in}}$ will maximise the covariance with the new test input $K(\mathbf{x}_{\text{in}}, \mathbf{x}_*)$. If there is no

**Algorithm 1** Hybrid scheme with on-the-fly machine learning

load initial database $C_p(X_p, X_p)$, $X_p = \{\mathbf{h}, \boldsymbol{\rho}\}_p$, $\langle \mathbf{p} \rangle$, $\boldsymbol{\theta}_p = \left\{\sigma_f^2, \ell^2, \sigma_n^2\right\}_p$
load initial database $C_q(X_q, X_q)$, $X_q = \{\mathbf{h}, \boldsymbol{\rho}, \mathbf{F}\}_q$, $\langle \mathbf{q} \rangle$, $\boldsymbol{\theta}_q = \left\{\sigma_f^2, \ell^2, \sigma_n^2\right\}_q$
set case variables $N, \mathbf{s}_{1\ldots N}$, $\mathbf{h}_{1\ldots N}$, $\boldsymbol{\rho}^1_{1\ldots N}$, $t_{\text{end}}$, $\Delta t$, $\sigma_t$
**procedure** TIME EVOLUTION
    **for** $i \leftarrow 1, t_{\text{end}}/\Delta t$ **do**
        $t \leftarrow i\Delta t$
        $X_* \leftarrow \{\mathbf{h}_{1\ldots N}, \boldsymbol{\rho}^i_{1\ldots N}\}$
        $\left\{\bar{\mathbf{p}}^i_{1\ldots N}, {\boldsymbol{\sigma}_p}^i_{1\ldots N}\right\} \leftarrow \text{GPREG}(C_p, X_p, \langle \mathbf{p} \rangle, X_*; \boldsymbol{\theta}_p)$ ▷ Eqs. (12), (13)
        $\boldsymbol{p\prime}^i_{1\ldots N} \leftarrow \Delta\bar{\mathbf{p}}^i_{1\ldots N}/\Delta\mathbf{S}_{1\ldots N}$
        set external forcing $F^i_{\text{ext}}$
        **for** $j \leftarrow 1, N$ **do**
            $F^i_j \leftarrow F^i_{\text{ext}} - (m/\rho^i_j)p\prime^i_j$ ▷ Eq. (4)
            $\mathbf{x}_* \leftarrow \left\{h^i_j, \rho^i_j, F^i_j\right\}$
            $\left\{\bar{q}^i_j, {\sigma_q}^i_j\right\} \leftarrow \text{GPREG}(C_q, X_q, \langle \mathbf{q} \rangle, \mathbf{x}_*; \boldsymbol{\theta}_q)$ ▷ Eqs. (12), (13)
            **if** ${\sigma_q}^i_j > \sigma_t$ **then**
                **procedure** PERFORM NEW MD SIMULATION
                    find nearest prior input $\mathbf{x}_{\text{in}}$
                    $t_s \leftarrow \text{PSEUDOMD}(\mathbf{x}_{\text{in}}, \mathbf{x}_*)$
                    $\left\{\bar{q}^i_j, {\sigma_q}^i_j, \bar{p}^i_j, {\sigma_q}^i_j\right\} \leftarrow \text{MD}(\mathbf{x}_{\text{in}}, \mathbf{x}_*, t_s)$
                    append $X_q \leftarrow [X_q\ \mathbf{x}_*]$
                    append $\langle \mathbf{q} \rangle \leftarrow \left[\langle \mathbf{q} \rangle\ \bar{q}^i_j\right]$
                    update $C_q \leftarrow \text{UPDATECOV}(C_q, X_q; \boldsymbol{\theta}_q)$
                    **if** $\left\{h^i_j, \rho^i_j\right\} \neq X_p$ **then**
                        append $X_p \leftarrow \left[X_p\ \left\{h^i_j, \rho^i_j\right\}\right]$
                        append $\langle \mathbf{p} \rangle \leftarrow \left[\langle \mathbf{p} \rangle\ \bar{p}^i_j\right]$
                        update $C_p \leftarrow \text{UPDATECOV}(C_p, X_p; \boldsymbol{\theta}_p)$
                    **end if**
                **end procedure**
            **end if**
        **end for**
        $\boldsymbol{\rho}^{i+1}_{1\ldots N} \leftarrow \Delta t(1/A)\Delta\bar{\mathbf{q}}^i_{1\ldots N}/\Delta\mathbf{S}_{1\ldots N}$ ▷ Eq. (3)
    **end for**
**end procedure**

database entry for the channel height to be tested, new MD simulations start with the fluid atoms in a lattice, with zero mean velocity. We estimate the time for the start-up phase $t_s$ by performing a simple 'pseudo MD' simulation using a 1D Navier-Stokes solver with a Navier slip condition.

After the simulation is complete, the measured mass flow rate $\bar{q}_j^i$ is appended to the vector $\langle \mathbf{q} \rangle$, the input $\mathbf{x}_*$ appended to the input matrix $X_q$, and the covariance matrix for the data $C_q([X_q\ \mathbf{x}_*], [X_q\ \mathbf{x}_*])$ is updated. If $\left\{h_j^i, \rho_j^i\right\}$ do not exist in the pressure database, then the pressure input matrix $X_p$, output vector $\langle \mathbf{p} \rangle$, and covariance matrix $C_p(\left[X_p\ \left\{h_j^i, \rho_j^i\right\}\right], \left[X_p\ \left\{h_j^i, \rho_j^i\right\}\right])$ are also updated.

## 3. Results and Discussion

Without recourse to experimental results, we test the accuracy of our hybrid scheme by comparing it to a full MD simulation for the same system; this also enables us to directly quantify the computational savings of our scheme. All the full MD solutions presented here are taken from Borg et. al. [28], with data points representing block averages over 2000 time-steps to reduce noise. For the majority of the results we present, the external forcing $F_{\text{ext}}$ is sinusoidal with an amplitude of $F_A = 0.487$ pN and a period of $T = 10.8$ ns (Case C in [28]), i.e.

$$F_{\text{ext}}(t) = F_A \sin\left(\frac{2\pi t}{T}\right). \tag{15}$$

This large oscillation period ensures a high degree of temporal scale separation. Let us first consider the case where we initially have an empty MD database. This could be considered the 'purest' form of machine learning in our scheme, as the input ranges that need to be learned are not influenced by human intuition. While it is fairly straightforward to estimate the input ranges for our benchmark system (as explained later), this may not be true for more complex systems, with a larger number of inputs; it is therefore important to demonstrate that our scheme is sufficiently robust to accurately model the flow behaviour, even when we have no prior information. Nevertheless, setting the target uncertainty threshold $\sigma_t$ involves some subjectivity. A sensible approach is for the uncertainty threshold to be set larger than the observation noise ($\sigma_{nq}$), because it is difficult for the model to make predic-

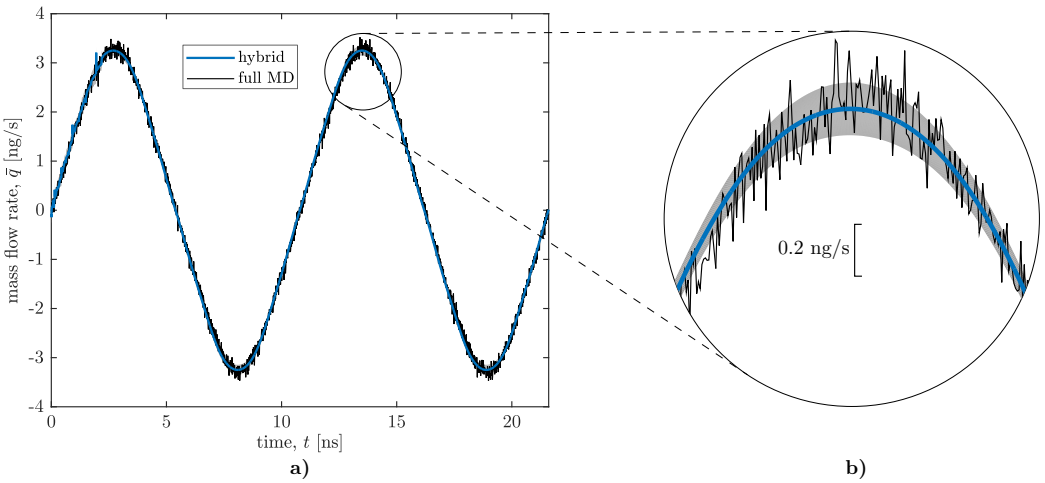


Figure 2: Transient mass flow rate results for our hybrid scheme (Case 1 — subdomain 1) and the full MD simulation [28], showing a) the full time series, and b) a close-up to highlight the uncertainty quantification of the hybrid solution — the grey envelope represents 1.96 standard deviations of the mean, representing the 95% confidence interval for our hybrid solution. The hybrid solution uses a tight uncertainty threshold of 0.1 ng/s and starts from an empty database.

tions with more accuracy than the data upon which it is based[4]. Noting from §2.3 that $\sigma_{nq} = 0.05$ ng/s, we initially choose a threshold of $\sigma_t = 0.1$ ng/s. This will be referred to as Case 1.

The transient mass flow rate results for Case 1 are displayed in Fig. 2, showing excellent agreement shown between the output of our hybrid scheme and the measurements from the full MD simulation. Mass conservation means that the mass flow rate profile is approximately the same at all subdomain locations, so we present the data only for subdomain 1. As we begin from an empty database, initially our hybrid scheme must perform MD simulations with high frequency, because there is limited data upon which to base a prediction. Therefore, the hybrid solution (blue line) exhibits noise similar to that of the full MD simulation up until $t = 2.7$ ns, where the external forcing function peaks. As the system geometry and external forcing function are both symmetric, after this time no 'new' input configurations are experienced, and no further MD simulations need to be performed. Beyond this

[4]If there are multiple data points within close proximity to the test input, then the uncertainty can be lower than the observation noise.

time, our hybrid solution near-perfectly captures the sinusoidal variation of mass flow rate, with smoothness resulting from our covariance prior. Figure 2b shows a close-up of the second mass flow rate peak and highlights that the uncertainty of our hybrid solution (grey region, representing 95% confidence bounds) is smaller than the noise in the full MD simulation. This is because the properties in the full MD simulation are transient over time-averaging for each data point, whereas our MD subdomain simulations are steady-state. The uncertainty is larger at the extremes of mass flow rate because these configurations exhibit the most extreme force and density inputs, so the model must extrapolate beyond its existing database.

While building the database during the first 2.7 ns, the model continually has to extrapolate for mass flow rate predictions (as the increasingly large forces and densities render existing data less relevant to the current conditions), until the uncertainty threshold is reached and a new MD simulation is performed. As the first mass flow rate peak approaches, there will be a final MD simulation, beyond which all configurations are sufficiently close to those in the database that the uncertainty of future predictions is always below the threshold. This means that at all subsequent peaks, the model must extrapolate from the database, whereas, for all other configurations, the model will interpolate within the database. One potential benefit of our machine-learning approach is the ease in which the effect that uncertainty in the external forcing function has on the mass flow rate solution can be quantified. Using pure MD simulations or traditional hybrid methods, it would be an extremely cumbersome task to model a wide range of forcing functions; however, with our scheme, once the database is built, new cases can be performed incredibly quickly (as shown in §3.3).

Figure 3 displays the transient density profiles for Case 1, calculated at each of the five subdomains. Our hybrid solution shows good agreement with the measurements from the full MD simulation, although it fails to produce the subtle oscillations in subdomain #1, and occasionally overestimates the density in subdomains #3 and #4. However, despite overestimating the density, our hybrid solution successfully captures the non-harmonic profile in these sections. The discrepancies our solution displays mirrors those in the solutions of Borg et. al., so are likely due to a common inaccuracy in the multiscaling, rather than an error in our GP regression model. We note that since density is an input to rather than an output of the regression model there is no simple method for uncertainty quantification.

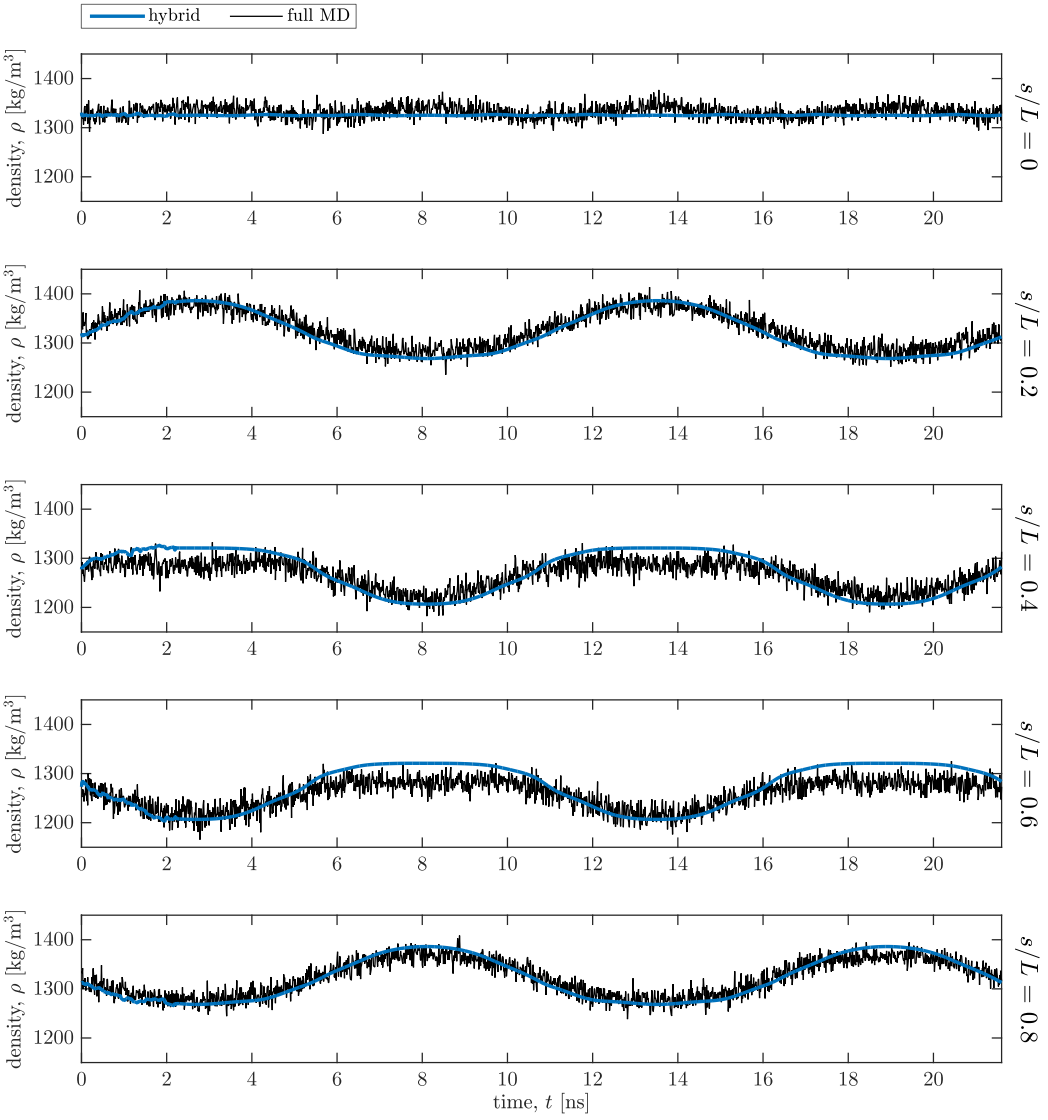


Figure 3: Transient density results for our hybrid scheme (Case 1) and the full MD simulation, calculated at each of the five subdomains. The hybrid solution uses a tight uncertainty threshold of 0.1 ng/s and starts from an empty database.

*3.1. The effect of the uncertainty threshold*

Cases T1 – T6 demonstrate the effect of the choice of uncertainty threshold on our hybrid solution. Case T1 is the same as Case 1, and the threshold increases up to 0.9 ng/s for Case T6. To isolate the effect of the threshold, Cases T1 – T6 all start with an empty database. Figure 4 confirms our intuition that the choice of threshold in our hybrid scheme is a trade-off between the accuracy and computational efficiency of its solution. As the threshold is raised, MD simulations are performed less frequently, so the accuracy of our hybrid solution drops, but the speed of computation increases. The signal standard deviation is $\sigma_{f_q} = 1$ ng/s (see §2.3), so when the threshold $\sigma_t > 1$ ng/s, MD simulations will never be performed, even when starting from an empty database. In this instance, the mass flow rate in each subdomain is predicted to be zero for the entire time series, as this is the assumed mean (see §2.2).

Figure 4a shows the total mass flow rate error $\epsilon$ for each case, relative to the error obtained if the database remains empty (which represents the maximum possible discrepancy), defined by

$$\epsilon_t = \frac{\sum_{i=1}^{t} |\bar{q}_i - Q_i|)}{\sum_{i=1}^{t} |Q_i|} \times 100\%, \tag{16}$$

where $\bar{q}_i$ is the mass flow rate, which is either measured by a MD subdomain simulation (when $\sigma_{q_i} > \sigma_t$), or predicted by our regression model (when $\sigma_{q_i} \leq \sigma_t$) at time index $i$ (averaged across all subdomains), and $Q_i$ is the mass flow rate from the full MD simulation. To obtain a meaningful error, the noise from the full MD solution is filtered by performing GP regression over the raw data, using a periodic kernel with time as the single input (see Appendix B for details). At first, the errors are greater than 100%. This is because the weak external forcing generates a near-zero mass flow rate, so the measurements from MD simulations (which are performed frequently at the start) are dominated by thermal fluctuations and masked by noise. Somewhat counterintuitively, the errors continue to drop after the final MD simulation is performed; this is because the weight of the errors incurred early in the trajectory diminishes as the hybrid solution becomes more accurate after the first peak in the forcing function. To obtain context for the error magnitudes, we also use equation (16) to calculate the 'error' in the raw full MD simulation data (grey area); for clarity, the calculation at time $t = 21.6$ ns is shown for the entire time series. Our results suggest that up to Case

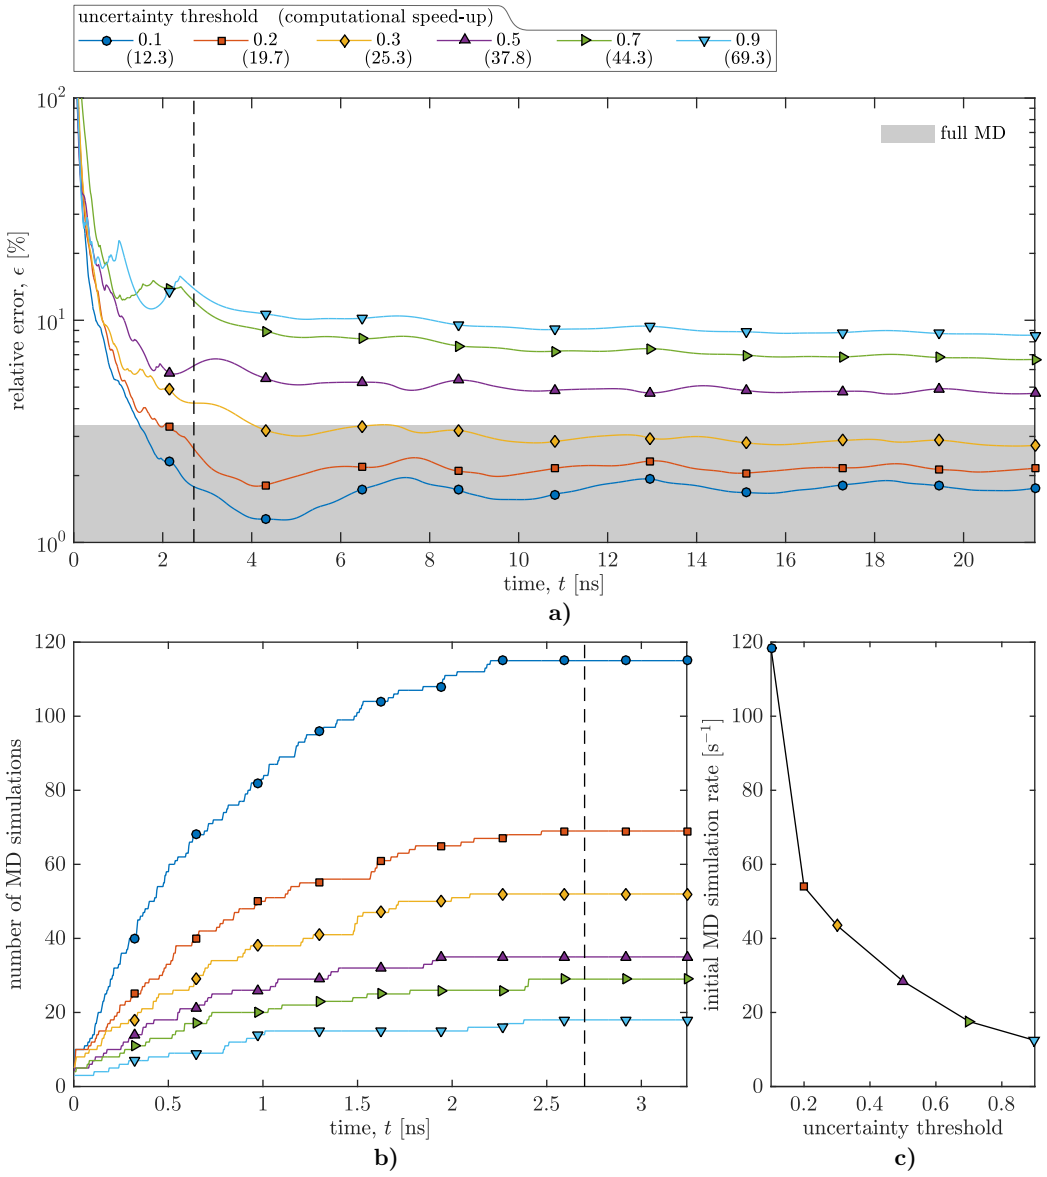


Figure 4: The influence of the uncertainty threshold $\sigma_t$ on the hybrid solution for mass flow rate (starting from an empty database): a) the transient relative error to the smooth full MD solution (the grey area shows the error due to noise in the full MD solution); b) the cumulative number of MD simulations; c) the initial MD simulation rate. The vertical dashed line in a) and b) denotes time at the first peak in the external forcing function, after which no MD simulations are performed. The data is from Cases T1 ($\sigma_t = 0.1$ ng/s) to T6 ($\sigma_t = 0.9$ ng/s), and the computational speed-up over the full MD solution for each case is displayed in parentheses in the legend.

T3, where $\sigma_t = 0.3$ ng/s, the error in our hybrid solution is comparable to the error in the full MD simulation due to noise.[5] Overall the error increases with increasing threshold, as is expected.

Figures 4b and c demonstrate how the computational speed of our scheme varies with the uncertainty threshold. Figure 4b shows the cumulative number of MD simulations, which is proportional to the efficiency of the solution. As expected, the trend is that the lower the threshold, the more MD simulations must be performed. In all cases no further MD simulations are required after $t = 2.7$ ns. The computational speed-up, given in brackets in the legend of Fig. 4, is calculated by

$$\text{speed-up} = \frac{t_{\text{end}} N_{a_f}}{\bar{t}_{\text{sim}} N_{\text{sim}} \bar{N}_{a_h}}, \tag{17}$$

where $N_{a_f}$ is the number of atoms in the full MD simulation, $\bar{t}_{\text{sim}}$ is the average time-steps performed in a single MD subdomain simulation, $N_{\text{sim}}$ is the number of MD subdomain simulations performed for the hybrid solution, and $\bar{N}_{a_h}$ is the average number of atoms in each of those MD simulations. For the tightest threshold (Case T1), our hybrid solution provides a modest speed-up over the full MD simulation of 12.3×; this rises to 69.3× for the loosest threshold (Case T6), however the relative error of this solution is 2.5× that of the full MD simulation. All cases (T1 – T6) all show logarithmic growth for the total number of MD simulations, with respect to time — i.e. the frequency of MD simulations decreases as the database becomes larger, and the predictions become more accurate.

Figure 4c shows how the initial required MD simulation rate varies with the uncertainty threshold. This can be considered a quantification of the 'novelty' of the solution — i.e. the rate of identifying new configurations that are dissimilar to those on which the model has been previously trained. We calculate this novelty by assuming the MD simulation rate is approximately constant over the first 0.27 ns.[6] The results confirm that lower thresholds produce a higher novelty as they require more certainty for a configuration to be recognised as being similar to those in the existing database.

[5] It would be wrong to say that our hybrid solution is more accurate than the full MD solution.

[6] For the variable threshold run, the $x-$axis in Fig. 4c is taken to be the average threshold over the first 0.27 ns.

*3.2. The effect of the initial database*

As previously acknowledged, prior to running our hybrid scheme, we can make some good estimations for the range of input values in our example case. For the particular system of interest studied here, we know that the channel height will always be equal to either $h_1$, $h_2$, or $h_3$ (see Fig. 1 for reference; it is likely that the magnitude of the total force $|F|$ will range from zero to at least the magnitude of the external force (0.487 pN); and from knowledge of nanofluidic MD simulations, we conservatively estimate that density will vary between 1120 kg/m$^3$ and 1480 kg/m$^3$ — approximately $\pm$180 kg/m$^3$ from the initial values. Simulations are performed at evenly-spaced increments across these ranges to create an initial database to train the model. Subsequently, 'new' configurations are encountered less frequently, and more predictions are made by interpolating between existing database entries, rather than extrapolating. This means predictions are made with more certainty, and the total number of simulations is likely to be reduced.

Table 1: Initial databases for Cases D1 – D5. See Fig. 1 for channel height references.

| Case | Channel heights | Initial database size |
|---|---|---|
| D1 | – | 0 |
| D2 | $h_1$ | 15 |
| D3 | $h_3$ | 16 |
| D4 | $h_1, h_3$ | 31 |
| D5 | $h_1, h_2, h_3$ | 47 |

The initial databases for cases D1 – D5 are outlined in Table 1, where Case D1 is the same as Case 1 — i.e. the initial database is empty. The other four databases are trained using one or more channel height and, for each height, four different densities and forces, evenly spaced across the estimated ranges above. The exception is for channel height $h_1$, the largest force (0.487 pN) applied to the lowest density fluid (1120 kg/m$^3$) produces a shear rate beyond the critical limit [29], so the mass flow rate does not converge; therefore, this data entry is removed from the initial databases. For all cases, the uncertainty threshold is set at $\sigma_t = 0.1$ ng/s.

The results for Cases D1 – D5 are presented in Fig. 5. Figure 5a demonstrates that the size of the initial database has a negligible effect on the overall accuracy of the method, with all cases exhibiting errors comparable to the noise in the full MD simulation. This result is expected, because the

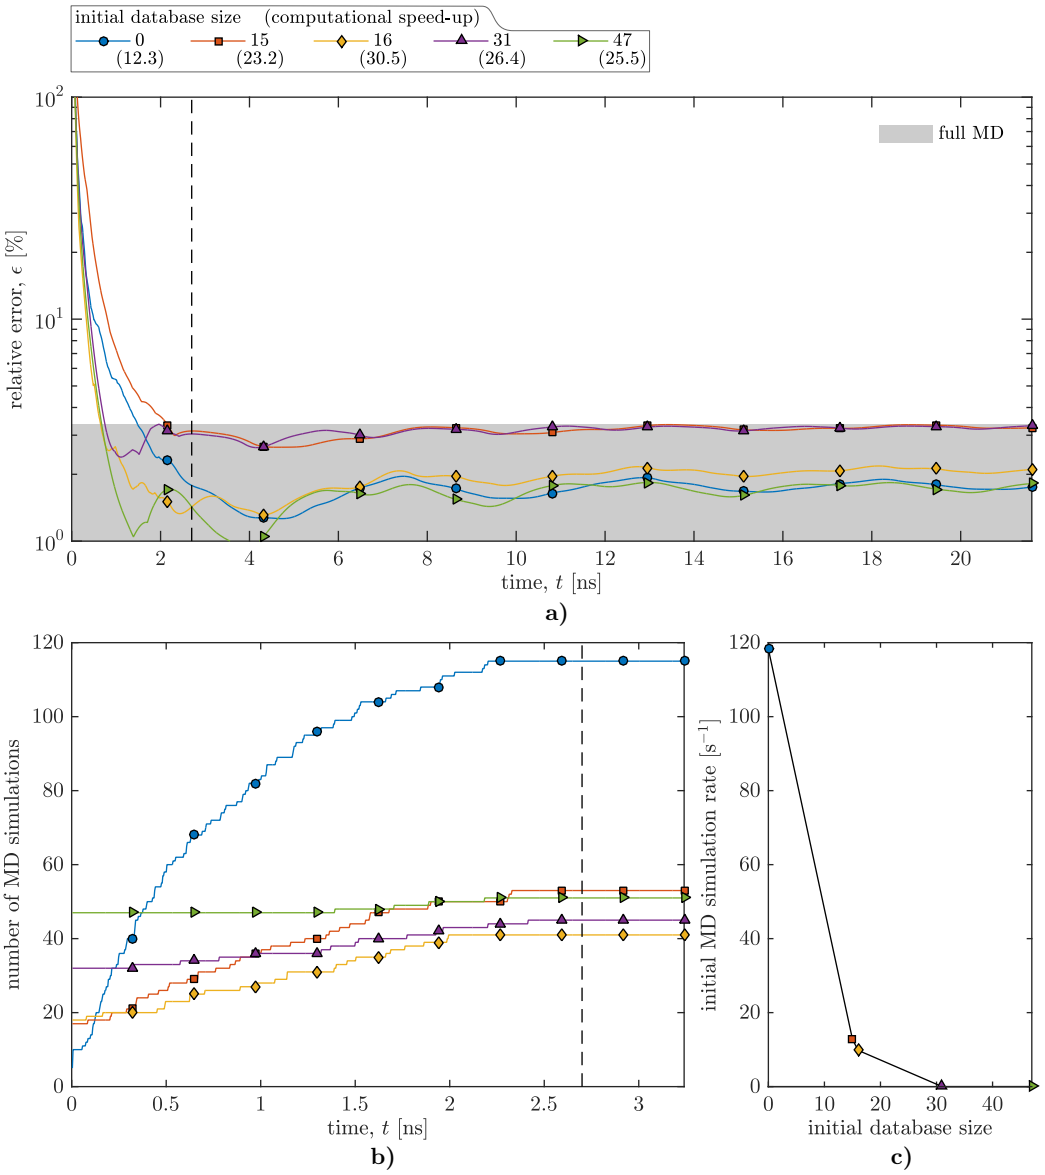


Figure 5: The influence of the initial database on the hybrid solution for mass flow rate (with an uncertainty threshold of 0.1 ng/s): a) the transient relative error to the smooth full MD solution (the grey area shows the error due to noise in the full MD solution); b) the cumulative number of MD simulations required; c) the initial MD simulation rate. The vertical dashed line in a) and b) denotes time for the first peak in the external forcing function, after which no MD simulations are performed. The data is from Cases D1 – D5 (see Table 1, and the computational speed-up for each case is displayed in parentheses in the legend.

uncertainty produced by extrapolating beyond a smaller initial database is countered by more frequently performing MD simulations.

Figure 5b shows that the total number of MD simulations decreases when the initial database is not empty, because more predictions are made through interpolation and 'new' configurations are not found so frequently at the start of the time series. This effect is demonstrated in Fig. 5c, where the novelty of the cases decreases as the initial database increases, reaching zero for the largest database in Case D5. However, the total number of MD simulations performed does not continue to fall as the initial database grows. For larger initial databases, sometimes redundant information is gathered which could have been inferred from neighbouring configurations; this excessive training outweighs the benefit of the increased confidence in subsequent predictions. This effect is demonstrated in the results for Case D5, where increasing the initial database to include simulations with channel height $h_2$ only negligibly reduces the number of 'on-the-fly' simulations performed compared to Case D4.

Another example of redundant information is the discrepancy between the results for Cases D2 and D3, Despite the model training on only one extra configuration for the latter case. This is due to the geometry of the case: a much larger force is required near the throat of the channel (subdomain #3) than at the inlet/outlet (subdomain #1) to generate equal mass flow rates.[7] As such, the local pressure gradient always acts in the opposite direction to the external force at the inlet/outlet of the channel, so the peak external force is never experienced in subdomain #1, and training on it for initial database provides little information. Conversely, all of the information is used when the initial database is formed using subdomain #3. Fewer MD simulations corresponds to an increase in computational efficiency — Case D3 is 30.5× faster than the full MD simulation while maintaining the same level of accuracy.

### *3.3. Building on an existing database*

As we have already demonstrated with previous results, one important advantage of our hybrid scheme is that it enables information to be reused so predictions can be made with more certainty from a continually-growing

[7]It is difficult to provide similar analysis for the density ranges, because at the nanoscale, viscosity varies rapidly with density, so greater density does not necessarily produce a larger mass flow rate.

database. So far, this information has been reused by increasing the confidence of our mass flow rate predictions for configurations later in the time series, within the same case. However, we can go further. For example, suppose having completed the hybrid simulation, we decided that we are really interested in a flow feature occurring at $s = 6.8$ nm (halfway between subdomains #1 and #2). Our previous options would have been to run the expensive full MD simulation to ensure every flow feature is captured, or to perform a new hybrid simulation using different subdomain locations; both of which are computationally wasteful. However, with our machine-learning-accelerated hybrid scheme, we can simply create a new case which has more subdomains, with the model already trained on the database that we generated in the previous case.

In Cases S1 – S4, we demonstrate how this approach can be used to continually add subdomains and refine the streamwise density profile. Case S1 is the same as case 1; in each subsequent case, the number of subdomains is doubled (with the new subdomain locations bisecting the old subdomain locations), and the initial database is that generated at the end of the previous case.[8] For all cases, the uncertainty threshold is $\sigma_t = 0.1$ ng/s. As the spacing between adjacent subdomains decreases, the relevance of data measured at neighbouring subdomains increases, and successively fewer MD simulations are performed, as shown in Fig. 6a. Using 40 subdomains, no new MD simulations are required at all. In addition, the streamwise density profiles calculated using five subdomains are inaccurate, because central differences are used to model spatial gradients, which assumes the variation between adjacent subdomains is linear. However, as the spacing between subdomains decreases, this linear assumption becomes more accurate. Figures 6b-d show how the streamwise density profiles for case S4 (40 subdomains) compares to the profile measured by the full MD simulation, at three snapshots in time. Our results show good qualitative agreement with the noisy MD data.

Another example of building on an existing database is evaluating the response to different external forcing functions in the same geometry. Again, without the aid of machine learning, this would require performing an entirely new hybrid or full MD simulation. Using the database generated at the end of Case 1 as our initial database, we perform two new cases: Case F1, with

[8]Note, extra pre-simulations for the pressure database are not required for the new subdomain heights, as the inference is already sufficiently accurate.

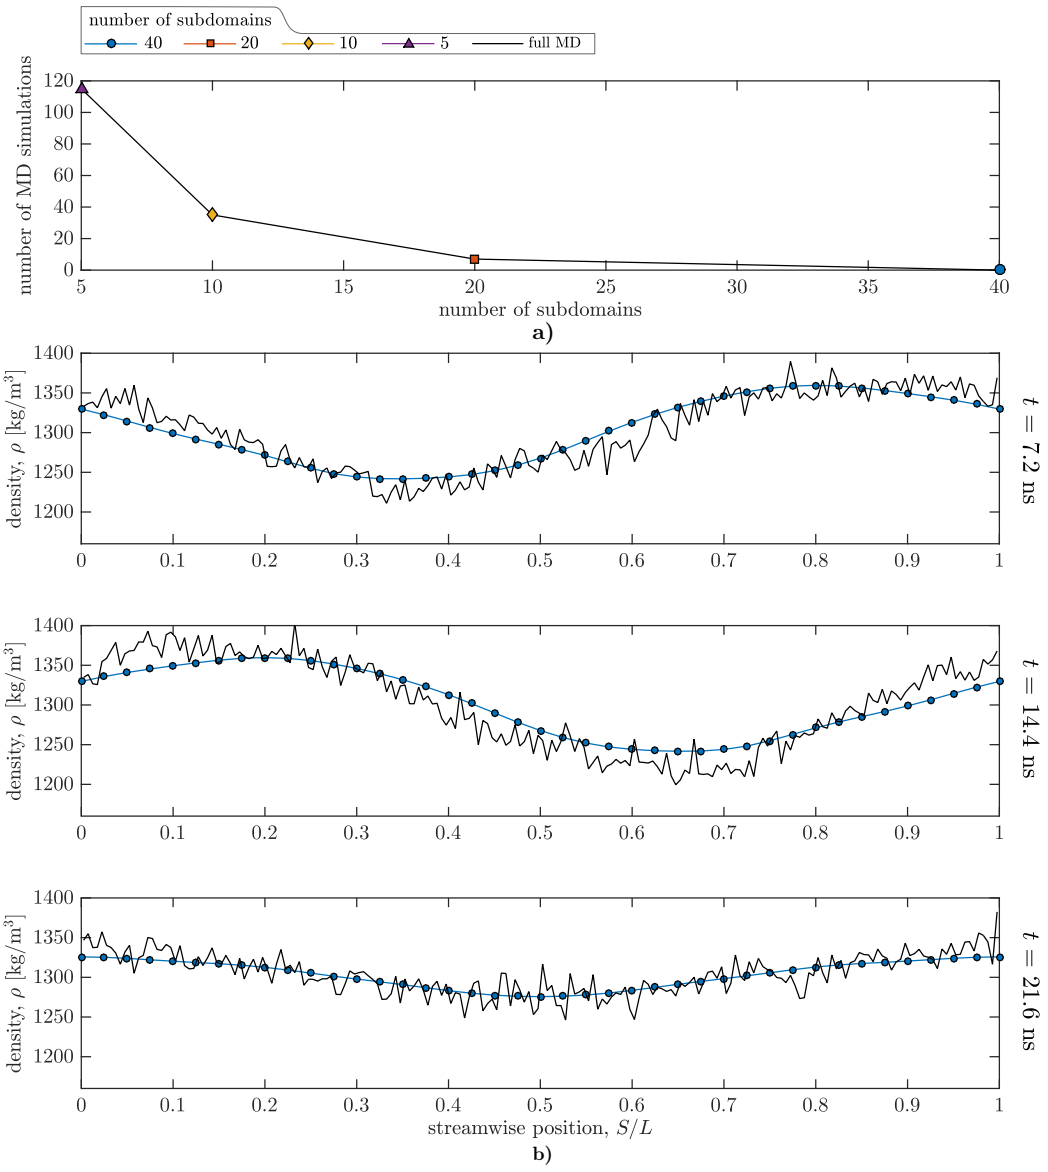


Figure 6: The influence of the number of subdomains on the hybrid solution for density (with an uncertainty threshold of 0.1 ng/s): a) the number of new MD simulations; and b) streamwise density profiles for the full MD simulation and the hybrid solution using 40 subdomains at $t = 7.2$ ns, $t = 14.4$ ns, and $t = 21.6$ ns. The data is from Cases S1 – S4.

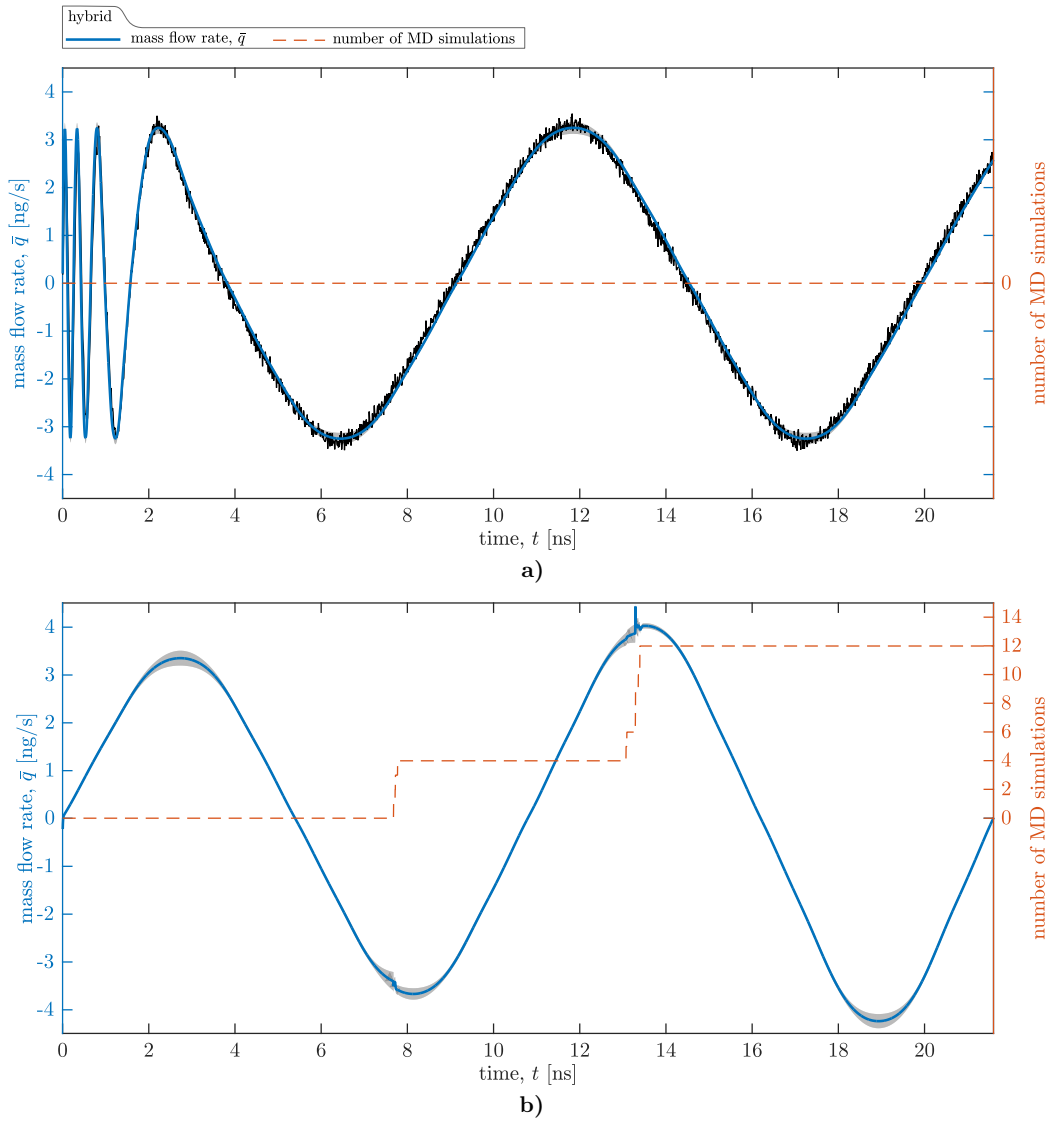


Figure 7: Transient mass flow rate results (left $y-$axis) for the hybrid solution of a) Case F1 (variable frequency external force), and b) Case F2 (variable amplitude external force). For Case F1, the result from the full MD simulation is also plotted. In both cases, the hybrid solution uses an uncertainty threshold of 0.1 ng/s, and an initial database generated at the end of Case 1. The cumulative number of new MD simulations is also plotted (right $y-$axis).

a variable-frequency external force; and Case F2, with a variable-amplitude external force. In Case F1, the oscillation period of the force starts from 0.22 ns and gradually increases to 10.8 ns. [28]; the amplitude is 0.487 pN as case 1. The left $y-$axis of Fig. 7a shows the transient mass flow rate results for our hybrid solution and the full MD solution (Case D in Borg et. al. [28]), and the right $y-$axis shows the number of new MD simulations performed. Once again, our solution exhibits strong qualitative agreement with the full MD simulation, and the computational speed-up is effectively infinite[9] as no new MD configurations are performed. In Case F2, the amplitude of the force starts from 0.487 pN and gradually increases to 0.609 pN;[10] the oscillation period is a constant 10.8 ns as in case 1. In Fig. 7b, we show our hybrid solution for the transient mass flow rate profile (left $y-$axis) and the number of new MD simulations performed (right $y-$axis). There is no full MD simulation data for this forcing function, so we present our results in part as a challenge to others. The results appear consistent with our previous findings: the mass flow rate profile peaks at the peaks of the forcing function, and the magnitude of the fourth peak is approximately $1.25\times$ the magnitude of the fourth peak in Case 1 — mirroring the variation of force amplitude. New MD simulations are only required as the mass flow rate (and therefore forcing function) approach peaks, where the GP model has to extrapolate beyond the existing force entries in the database. The lack of smoothness in our hybrid solution, where new MD simulations are required, can be remedied by re-running the case with the final frozen database, at a negligible cost.

## 4. Conclusion

We have presented an enhancement to existing hybrid methods in fluid dynamics that uses on-the-fly machine learning techniques to predict the output of new system configurations, without having to continually perform atomistic simulations. This procedure enables information to be reused multiple times, drastically increasing the computational efficiency, while also providing the potential for simple uncertainty quantification.

[9] Of course, the speed-up is not actually infinite, but the computational effort to run the surrounding algorithm is neglible compared to an MD simulation.

[10] A much large force would cause the critical shear limit to be exceeded.

We compare our new scheme to full molecular dynamics (MD) simulations and find strong agreement, with errors within the range of thermal noise when a tight uncertainty threshold is set. As this threshold is raised, the error increases to up to 2.5× the MD simulation noise; however, the computational speed-up over a full MD simulation also increases. When starting from an empty database, raising the threshold from 0.1 ng/s to 0.9 ng/s increases speed-up from 12.3× to 69.3×. Thus, the choice of threshold is a trade-off between the required accuracy and computational efficiency.

We demonstrate the computational benefit of creating an initial database to train our predictive model, by estimating the expected range of inputs. This enables more predictions to be made via interpolation between data points, which provides less uncertainty than extrapolation leading to a requirement for fewer MD simulations to be performed 'on-the-fly'. While maintaining the same level of accuracy, starting with a modest initial database covering just a few data points can result in a speed-up of 30.5× the full MD simulation. Finally, we show how existing databases can be built upon (while never needing to be fully complete) to rapidly obtain high resolution hybrid solutions — i.e. cheaply add more subdomains at locations of interest — or to model different flow fields effectively instantly — i.e. no new MD simulations are required.

**Acknowledgements**

We would like to thank Matthew Borg for providing the data for the full MD simulations. This work is financially supported in the UK by EPSRC Programme Grants EP/I011927/1 and EP/N016602/1, and EPSRC grants EP/K038664/1 and EP/L027682/1. The computing facilities were provided by the Centre for Scientific Computing of the University of Warwick with support from the Science Research Investment Fund.

## Appendix A. MD parameters

Non-equilibrium MD simulations of dense fluid argon are performed using mdFoam [30, 31] — an in-house solver developed in OpenFOAM. Atoms are modelled as hard spheres which interact using shifted Lennard-Jones (LJ) pair potentials, with wall atoms frozen in place. Our MD parameters are necessarily identical to those used by Borg et. al. [28] for the full MD

simulations. The LJ potential

$$U_{LJ}(r_{ij}) = 4\epsilon \left[ \left( \frac{\sigma}{r_{ij}} \right)^{12} - \left( \frac{\sigma}{r_{ij}} \right)^{6} \right], \tag{A.1}$$

has a cut-off radius of 1.36 nm; for the fluid-fluid interactions, the LJ characteristic length and energy are $\sigma_{f-f} = 0.34$ nm and $\epsilon_{f-f} = 1.65678 \times 10^{-21}$ J, respectively; for the wall-fluid interactions, these parameters are $\sigma_{w-f} = 0.255$ nm and $\epsilon_{w-f} = 0.33 \times 10^{-21}$ J, respectively. The mass density of the wall atoms is $6.809 \times 10^3$ kg/m$^3$, and the mass of a single atom is $m = 6.6904 \times 10^{-26}$ kg. Fluid atom dynamics are described by Newton's laws of motion, which are numerically integrated using the Velocity Verlet method [32], with a time-step of 5.4 fs. The excess heat generated by applying an external force is removed by modifying the velocities in the $z$-direction (see Fig. 1) using a Berendsen thermostat; this ensures a thermally homogeneous system, maintained at a temperature of 292.8 K. All of our simulations are periodic in the $s$- and $z$-directions.

## Appendix B. Gaussian process regression of the full MD simulation data

To obtain error estimates for our hybrid solution, the full MD mass flow rate profile is modelled using a Gaussian process with a single input: time $t$. As the underlying profile is sinusoidal, we use a periodic kernel, i.e.

$$K(t, t') = \sigma_f^2 \exp\left( -\frac{2\sin^2\left(\pi|t - t'|/T\right)}{\ell^2} \right), \tag{B.1}$$

where $T = 10.8$ ns is the period of oscillation; the covariance function is still calculated using equation (10). For consistency with our hybrid scheme, we choose the signal standard deviation $\sigma_f = 1$ ng/s. We anticipate the noise standard deviation to be larger than in our quasi-steady subdomain simulations, due to transience of properties during the averaging period in the full MD simulation. Therefore, we optimise for the length scale and the noise standard deviation by maximising the log marginal likelihood of the raw full MD data (up to the first peak), as outlined in §2.3, and find that $\ell = 1$ and $\sigma_n = 0.1$ ng/s (twice that of our MD subdomain simulations). Gaussian process regression is performed over the entire time series using equations (12) and (13), with the mean mass flow rate assumed to be zero.

## References

## References

[1] R. Schoch, J. Han, P. Renaud, Transport phenomena in nanofluidics, Rev. Mod. Phys. 80 (3) (2008) 839–883.

[2] N. G. Hadjiconstantinou, Hybrid atomistic-continuum formulations and the moving contact-line problem, J. Comput. Phys. 154 (2) (1999) 245–265.

[3] M. Brenner, X. Shi, S. Nagel, Iterated instabilities during droplet fission, Phys. Rev. Lett. 73 (25) (1994) 3391–3394.

[4] G. Karniadakis, A. Beskok, N. Aluru, Microflows and Nanoflows: Fundamentals and Simulation, Springer, 2005.

[5] A. Patronis, D. A. Lockerby, Multiscale simulation of non-isothermal microchannel gas flows, J. Comput. Phys. 270 (2014) 532–543.

[6] K. Ritos, M. K. Borg, D. A. Lockerby, D. R. Emerson, J. M. Reese, Hybrid molecularcontinuum simulations of water flow through carbon nanotube membranes of realistic thickness, Microfluid. Nanofluid. 19 (2015) 997–1010.

[7] M. P. Allen, D. J. Tildesley, Computer Simulation of Liquids, Clarendon Press, 1987.

[8] G. A. Bird, Molecular Gas Dynamics and the Direct Simulation of Gas Flows, Oxford Engineering Science Series 42, Clarendon Press, 1994.

[9] H. S. Wijesinghe, N. G. Hadjiconstantinou, Discussion of hybrid atomistic-continuum methods for multiscale hydrodynamics, Int. J. Multiscale Comput. Eng. 2 (2) (2004) 189–202.

[10] N. G. Hadjiconstantinou, Discussion of recent developments in hybrid atomistic-continuum methods for multiscale hydrodynamics., Bull. Pol. Acad. Sci.: Tech. Sci. 53 (4) (2005) 335–342.

[11] M. Kalweit, D. Drikakis, Multiscale methods for micro/nano flows and materials, J. Comput. Theor. Nanosci. 5 (9) (2008) 1923–1938.

[12] K. M. Mohamed, A. A. Mohamad, A review of the development of hybrid atomistic-continuum methods for dense fluids, Microfluid. Nanofluid. 8 (3) (2009) 283–302.

[13] X. Nie, S. Chen, M. O. Robbins, Hybrid continuum-atomistic simulation of singular corner flow, Phys. Fluids 16 (10) (2004) 3579.

[14] M. K. Borg, D. A. Lockerby, J. M. Reese, Fluid simulations with atomistic resolution: a hybrid multiscale method with field-wise coupling, J. Comput. Phys. 255 (2013) 149–165.

[15] A. Patronis, D. A. Lockerby, M. K. Borg, J. M. Reese, Hybrid continuum-molecular modelling of multiscale internal gas flows, J. Comput. Phys. 255 (2013) 558–571.

[16] D. J. C. MacKay, Information theory, inference and learning algorithms, Cambridge Univ Press, Cambridge, 2003.

[17] M. C. Kennedy, A. O'Hagan, Bayesian calibration of computer models, J. R. Stat. Soc. Series B Stat. Methodol. 63 (3) (2001) 425–464.

[18] C. E. Rasmussen, C. K. I. Williams, Gaussian processes for machine learning., MIT Press, 2006.

[19] M. Aldegunde, J. R. Kermode, N. Zabaras, Development of an exchange–correlation functional with uncertainty quantification capabilities for density functional theory, J. Comput. Phys. 311 (2016) 173–195.

[20] P. Angelikopoulos, C. Papadimitriou, P. Koumoutsakos, Data driven, predictive molecular dynamics for nanoscale flow simulations under uncertainty, J. Phys. Chem. B 117 (47) (2013) 14808–14816.

[21] J. Behler, M. Parrinello, Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces, Phys. Rev. Lett. 98 (14) (2007) 146401.

[22] A. P. Bartók, M. C. Payne, R. Kondor, G. Csányi, Gaussian Approximation Potentials: The Accuracy of Quantum Mechanics, without the Electrons, Phys. Rev. Lett. 104 (13) (2010) 136403.

[23] W. J. Szlachta, A. P. Bartók, G. Csányi, Accuracy and transferability of gaussian approximation potential models for tungsten, Phys. Rev. B Condens. Matter 90 (10) (2014) 104108.

[24] Z. Li, J. R. Kermode, A. De Vita, Molecular dynamics with On-the-Fly machine learning of Quantum-Mechanical forces, Phys. Rev. Lett. 114 (March) (2015) 096405.

[25] V. Botu, R. Ramprasad, Adaptive machine learning framework to accelerate ab initio molecular dynamics, Int. J. Quantum Chem. 115 (16) (2015) 1074–1083.

[26] M. Caccin, Z. Li, J. R. Kermode, A. De Vita, A framework for machine-learning-augmented multiscale atomistic simulations on parallel supercomputers, Int. J. Quantum Chem. 115 (16) (2015) 1129–1139.

[27] N. Asproulis, D. Drikakis, An artificial neural network-based multiscale method for hybrid atomistic-continuum simulations, Microfluid. Nanofluid. 15 (4) (2013) 559–574.

[28] M. K. Borg, D. A. Lockerby, J. M. Reese, A hybrid molecular-continuum method for unsteady compressible multiscale flows, J. Fluid Mech. 768 (2015) 388–414.

[29] P. A. Thompson, S. M. Troian, A general boundary condition for liquid flow at solid surfaces, Nature 389 (6649) (1997) 360–362.

[30] G. B. Macpherson, J. M. Reese, Molecular dynamics in arbitrary geometries: parallel evaluation of pair forces, Mol. Simul. 34 (1) (2008) 97–115.

[31] M. K. Borg, G. B. Macpherson, J. M. Reese, Controllers for imposing continuum-to-molecular boundary conditions in arbitrary fluid flow geometries, Mol. Simul. 36 (10) (2010) 745–757.

[32] W. C. Swope, H. C. Anderson, P. H. Berens, K. R. Wilson, A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: Application to small water clusters, J. Chem. Phys. 76 (1) (1982) 637–649.